



Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

Fujitsu's ETERNUS DX400 S2 midrange storage arrays are the successor to their earlier ETERNUS4000 and DX400 products. The ETERNUS DX400 S2 products are sold worldwide via Fujitsu companies and channel partners.

Products and vendors that compete with the DX440 S2 and DX410 S2 systems include VNX and CX4 series from EMC, the AMS series from Hitachi Data Systems (HDS), the EVA from Hewlett Packard (HP), IBM DS5000 and DS4000 subsystems, and NetApp FAS 3200 series systems.

HIGHLIGHTS

- Multiple connection options
 - Up to 32 FC host ports
 - Up to 16 iSCSI 1/10 Gb Ethernet ports
- Resource sharing and access security with a LUN masking implementation called Host Affinity
- Support for RAID0, 1, 5, 5+0, 6 and RAID1+0 with global hot spare disks
- Online logical volume expansion
- RAID group sizes of 4 to 32 disk drives for RAID1+0 and 3 to 16 for RAID5.
- Block level data protection with error detection codes added to every data block
- Duplexed write cache with ECC protection
- Write cache data written to non-volatile storage if power fails
- Redundant paths from controllers to each disk
- Supports intermix of SAS, Nearline SAS, & SSD drives within same disk enclosure
- Call home feature for proactive maintenance
- Up to 96GB of cache
- Supports up to 960 disk drives
- Maximum raw capacity of 960TB
- Support for 16,384 logical volumes
- Optional Data protection software
 - Snapshot point-in-time copy
 - Equivalent Copy creates clones and mirrors.
 - Support for remote copy in synchronous or asynchronous mode
 - Remote copy using fibre channel and iSCSI
- Disk Traffic Control (DTC) performance feature to separate I/O based on sequential or random access.
- ETERNUS ECO-Mode for reduced power for infrequently accessed data (MAID)
- Storage controller based Disk Encryption to protect data at rest
- WAN Optimization and encryption of data using iSCSI remote adapters

Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

OVERVIEW

There are two models of the ETERNUS DX400 S2 that vary in capacity, performance, and connectivity. Heterogeneous server attachment is supported (individual server operating system support is listed in the Connectivity section of this document).

The ETERNUS DX400 S2 is part of the ETERNUS product line featuring the entry level ETERNUS DX60/80/90, and the high-end ETERNUS DX8000, that also provides mainframe connectivity in the Japanese market.

As seen below in Figure 1 that depicts the two models, either system may be rack mounted or configured in Fujitsu racks. Additionally, a minimum configuration consists of the one controller enclosure containing two controller modules, along with one or more disk tray enclosures.



Figure 1: ETERNUS DX400 S2 Product Line—Model DX410 S2, Model DX440 S2

Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

Model Comparison

		Model DX410 S2	Model DX440 S2
Form Factor		Rack Mount (3U base + 2U Expansion)	Rack Mount (3U base + 2U Expansion)
Number of Controllers		2	2
RAID Levels		0, 1, 1+0, 5,5+0, 6	0, 1, 1+0, 5,5+0, 6
Maximum Disk Drives		240 (3.5") / 480 (2.5")	480 (3.5"), 960 (2.5")
Disk Drive Types	3.5" or 2.5" Enterprise SSD	100 GB, 200 GB	100 GB, 200 GB
	3.5" 15k rpm SAS	300, 450, 600 GB	300, 450, 600 GB
	2.5" 10k rpm SAS	300, 450, 600 GB	300, 450, 600 GB
	3.5" 7.2k SAS nearline	1 TB, 2 TB	1 TB, 2 TB
Total Physical Capacity	Enterprise SSD	96.0 TB	192.0 TB
	SAS	288.0 TB	576.0 TB
	Nearline SAS	480.0 TB	960.0 TB
Maximum Volumes / RAID group		128	128
Max Number of Connected Hosts		1,024	1,024
Max Affinity Groups		1,024	1,024
Max Hosts per port		256	256
Thin Provisioning Pools		240	256
Maximum Logical Volumes		16,384	16,384
Maximum LUN Size		128 TB	128 TB
Number of Host Interfaces		4 - 16 (2 – 8 / controller)	4 - 32 (2 - 16 / controller)
Host Interfaces / Controller	FC	4/8/16 @ 2 /4 /8 Gb	4 - 32 @ 2 /4 /8 Gb
	FCoE	4/8 @ 10 Gb	4 - 16 @ 10 Gb
	iSCSI	4/8 @ 1 /10 Gb	4 - 16 @ 1 /10 Gb
Cache Capacity		8 or 16 GB / controller	6 - 48 GB / controller
Drive Interface		SAS (4X @ 6 Gbit/s)	SAS (4X @ 6 Gbit/s)
Dimensions (W x D x H)		19" x 27.6" x 5.2" (3U) Minimum Config	19" x 27.6" x 5.2" (3U) Minimum Config
Additional Drive Enclosures		Max 20	Max 40

¹Calculated as 1k byte = 1,000 bytes and includes system disk capacity

²Calculated as 1k byte = 1,024 bytes and when formatted as RAID5. User capacity depends on system environment

Table 1: Model Comparison



Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

PRODUCT ARCHITECTURE

The ETERNUS DX400 S2 is designed for the midrange storage market with two models available, the DX410 S2 and DX440 S2. Both models are available as rack mount or mounted in Fujitsu racks and support a wide range of configuration options to meet performance, capacity and connectivity requirements.

The base unit for the ETERNUS DX410 S2 contains one controller enclosure and up to eight (8) drives enclosures, which will house up to 120 disk drives. The expansion rack contains from 2 to 10 drive enclosures and will house up to 150 disk drives. A maximum configuration is one expansion rack for the Model DX410 S2 and two expansion racks for the Model DX440 S2. The maximum number of drives for Model DX410 S2 is 480 and for the Model DX440 S2 is 960 disk drives.

The drive interface between the controller and the disk shelves is 6 Gb SAS, using pairs (1 for each CM) of quad port, QSFP connectors.

Figures 2 and 3 illustrate the Models 410 and 440 along with their respective topologies.

Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

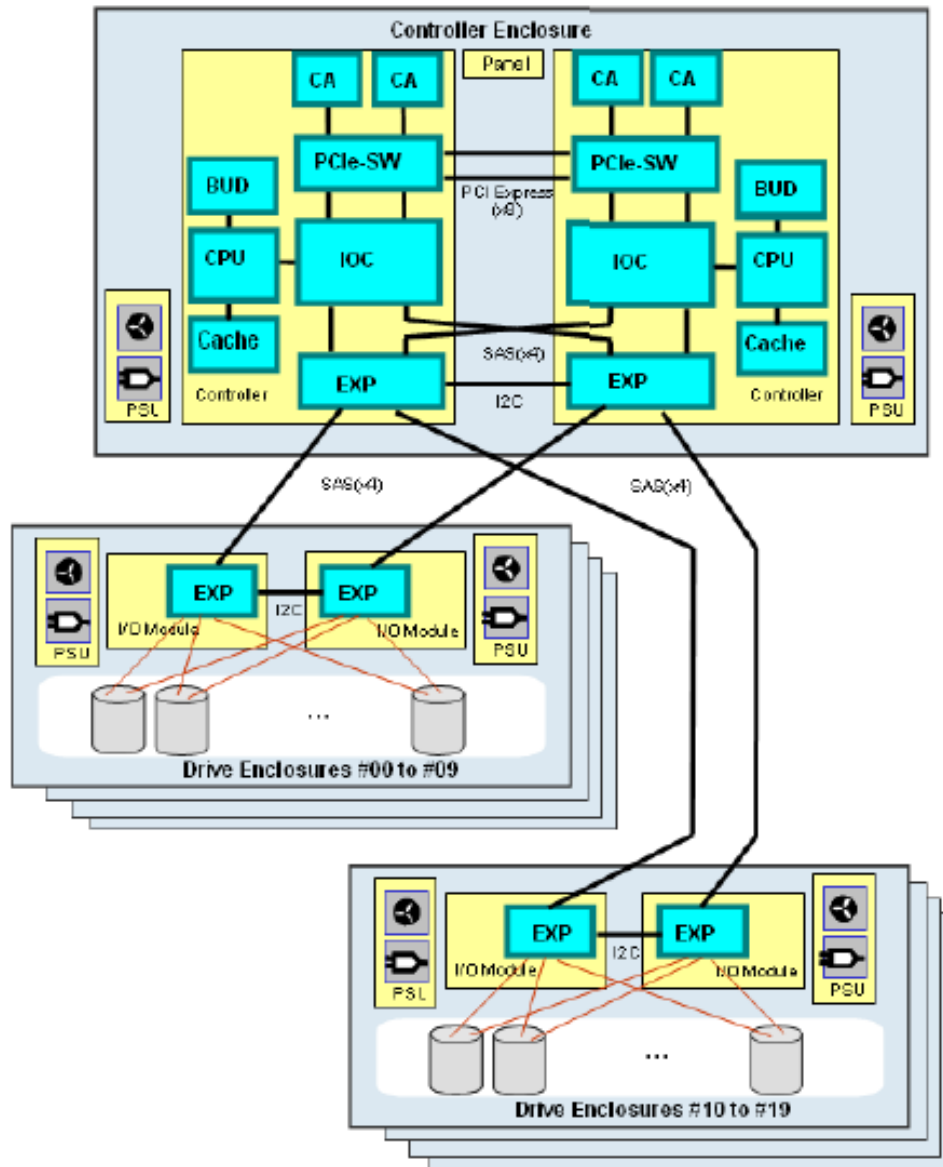


Figure 2: ETERNUS DX410 S2

Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

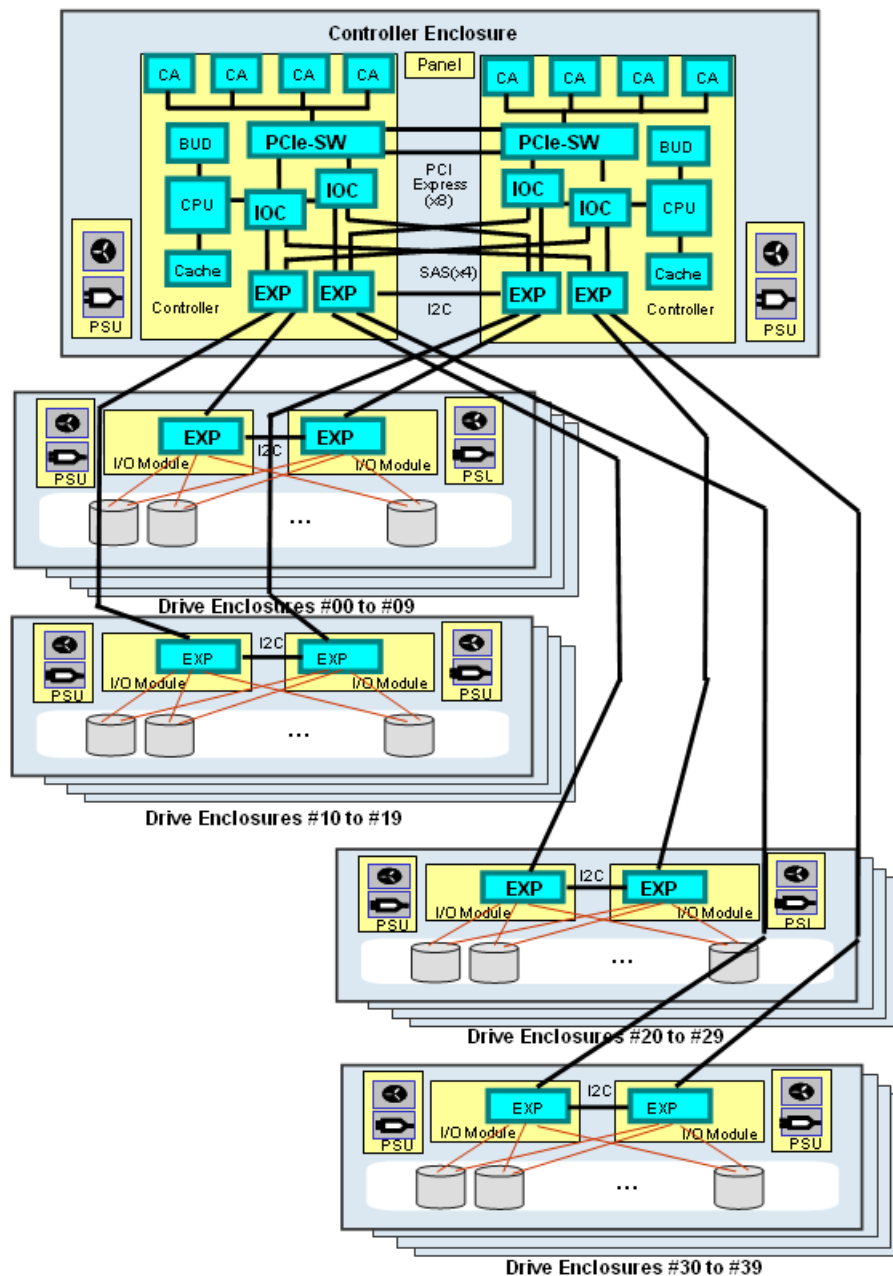


Figure 3: ETERNUS DX440 S2

Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

Controllers

Both models are dual controller and use a shared PCI-Express bus architecture for communication between the two controllers and for the attachment of the channel adapters (to the host) and device adapters (to the drives). The controllers execute the software that runs the storage systems. Controllers can run active-passive, or in parallel (active-active) for faster performance. The device adapters, which attach the drives to the controller, are attached with SAS Expanders supporting 6Gbps per path.

The ETERNUS DX410 S2 has a single two-core processor per controller, and the DX440- S2 has a single four-core processor. On the DX410, one core is assigned for processing, while the other core is used to perform encryption and parity creation for RAID5 and RAID6. In the DX440, two cores are assigned to processing and two cores to RAID and encryption calculations.

Host interface connectivity (Channel Adapters) for the ETERNUS DX400 S2 can be fibre channel, iSCSI, or FCoE or a mixture. The Model DX440 S2 supports a total of 96GB of cache in 6GB increments, with the Model DX410 S2 supporting 8 or 16GB. All configurations utilize cache split across two controller cards.

Cache is utilized for multiple purposes, including for read and write operations, as well as maintaining information required for data protection and replication functions such as Fujitsu's array based replication and point in time copy data protection software. Cache is allocated dynamically, with more detail covered on page 9 under the section "Cache" in this document.

Channel Adapters

The Channel Adapters (CAs) provide fibre channel, iSCSI, and FCoE interfaces to the host (server) system. They process commands from the host and manage accesses to the cache. The channel adapters plug into a PCI-Express bus on the controller module. Up to eight fibre channel host connections per controller may be installed (two quad channel adapter cards). For service, an entire controller module must be taken out of operation (a failover to another controller module must be done) before the channel adapter is replaced. After replacement, the controller module can be returned to service.

The channel adapters support fibre channel connectivity at 2Gbps, 4Gbps, and 8Gbps. Channel adapters support both 1Gbps and 10Gbps iSCSI, as well as 10Gbps FCoE

Table 2 below summarizes channel adapter connectivity.

Interface	Model DX410 S2	Model DX440 S2
Number of FC (8Gb) ports	4 - 8	4 - 32
Number of iSCSI (1Gb) ports	4 - 8	4 - 16
Number of iSCSI (10Gb) ports	4 - 8	4 - 16
Number of FCoE (10Gb) ports	4 - 8	4 - 16

Table 2: ETERNUS DX400 S2 Models Host Connectivity

Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

Device Interfaces

Integrated with the controller module, the Device Interface adapters (DIs) connect the disk drives to the controller modules using a 6 Gb SAS connection. The Model DX410 S2 device adapter has two SAS connections per controller; the Model DX440 S2 has four SAS connections per controller.

Drive Enclosures

The Fujitsu ETERNUS DX410 S2 and DX440 S2 systems support two types of enclosures:

- 2.5" drive enclosure, 24 slot – 2 I/O modules
- 3
- 3.5" drive enclosure, 12 slot – 2 I/O modules

RAID Groups

Fujitsu supports several RAID group configurations for both traditional and SSD disk devices. The configurations are shown below in Table 2. The maximum number of RAID groups is dependent upon the number of disk drives in the system, and the configuration guidelines provided below.

The Fujitsu ETERNUS DX systems utilize RAID groups as the first unit of grouping and data protection. Host accessible volumes (or LUN's) are created by aggregating or splitting up RAID groups. Logical volumes are made available to servers as LUNs (Logical Unit Numbers).

The ETERNUS DX systems support multiple RAID levels and disk combinations, providing different levels of protection and performance. The RAID levels supported include: RAID0, RAID1, RAID1+0 (data is mirrored on multiple disks and then striped across 4 to 32 disks), RAID5 with rotating distributed parity on multiple disks, RAID5+0 (which combines block striping of RAID0 with RAID5 rotating parity) and RAID6 which offers dual parity per RAID group.

Listed below in Table 2 are the allowable and recommended RAID configurations.

RAID	Drives / RAID Group	Recommended Drives Configuration (Grouping)
RAID0	2 – 16	Not recommended due to no protection!
RAID1	2	2(1D+1M)
RAID1+0	4 – 32	4(2D+2M), 6(3D+3M), 8(4D+4M), 10(5D+5M)
RAID5	3 – 16	3(2D+1P), 4(3D+1P), 5(4D+1P), 6(5D+1P)
RAID5+0	6 – 32	6[(2D+1P)x2], 8[(3D+1P)x2], 10[(4D+1P)x2], 12[(5D+1P)x2]
RAID6	5 – 16	5(3D+2P), 6(4D+2P), 7(5D+2P),

Table 3: Supported RAID Groups and Recommended Configurations

With the release of the "S2" update to the DX410 and DX440, Fujitsu has removed earlier restrictions on the number of SSD drives supported. A full configuration of SSD drives is now supported, along with the RAID group configurations outlined above in Table 2.



Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

Hot Spare

Hot spare disk drives are global spares that may be used by any RAID group with access to the hot spare. The hot spares allow rebuilding to start automatically when a disk drive fails, or before failure if a high drive error rate is detected. After replacing a failed drive, the data is copied to the replaced drive and the original spare returned to spare status. If using different capacity drives in the system, the hot spare must be equal to or greater than the largest capacity drive. Hot spares are defined to the system via the ETERNUSmgr function.

When SAS, nearline SAS, and SSD drives are installed, hot spare disks must be defined for each type.

When invoking hot sparing, a hot spare disk with the same capacity as the failed disk will be used first; if one does not exist, the largest capacity hot spare disk will be used. There is no limit for the number of hot spares that can be defined.

Disk Drives

The ETERNUS DX410 S2 and DX440 S2 support multiple drive types, capacities and physical sizes. The disk drive types correspond to different workloads, IO centric for SSD, mixed performance capacity for enterprise 15k and 10k rpm drives, and capacity centric for the nearline 7.2k rpm drives. Both SFF (2.5") and LFF (3.5") form factor drives are supported on the ETERNUS DX410 S2 and DX440 S2.

All drives utilize a Serial Attached SCSI (SAS) 6 Gbit/s interface with dual ports for high availability.

Evaluator Group Comment: Fujitsu has chosen to support both 3.5" and 2.5" drives, providing significant choice for customers as the industry transitions to the newer 2.5" drive form factor. Additionally, the ETERNUS DX410 S2 and ETERNUS DX440 S2 systems support SSD drives in both small and large form factor. Fujitsu is one of the few vendors offering SSD drives in this price class, which provides users with more choices when configuring for performance and capacity centric applications.

MAID - ECO-Mode

Fujitsu utilizes Massive Array of Idle Disks (MAID) technology for infrequently accessed data in order to reduce power and cooling requirements. This feature is referred to as 'ETERNUS ECO-Mode'. By allowing RAID groups that are infrequently accessed to spin down when they are not used, the total power consumption may be reduced.

The ECO-Mode feature controls the disk drive rotation within a RAID group, affecting all drives that form a RAID group. This feature is managed through ETERNUSmgr. ECO Mode is disabled when system activity requires access to the RAID group, including for data protection operations, mirror copying, RAID group rebuild operations and other system required activity.

ECO Mode utilizes a setting, in conjunction with RAID group activity to determine if a RAID group should be spun down. If there is no activity, and the setting indicates the group should be spun down, then ECO Mode takes effect. If either of these two conditions is not met, then the group will remain spinning. ECO Mode utilizes time frames for a specific RAID group. RAID groups may be enabled for the following time frames:

Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

- **Everyday** - The mode is enabled every day from the specified start time to the specified end time.
- **Every week** - The mode is enabled every day from the specified start time to the specified end time.
- **Specific day** – The day or date of the month, along with start, and number of days is specified
- **Specific week** – A day of the week within the month, along with the start and ending day of the week and the starting and ending time are specified
- **Last Access Time** – A specific amount of time has elapsed since a RAID group was last accessed, this may be set to 10, 20, 30, 40, 50 or 60 minutes (default is 30 minutes)

Evaluator Group Comments: *The lower power consumption of ECO-Mode is a somewhat unique feature in the midrange market segment that provides customers with greater flexibility in matching their storage systems to business requirements.*

Disk Data Encryption

ETERNUS provides data encryption and protects critical, security sensitive data from unauthorized access to drives that have been removed from the storage system. Setup and management of encryption is via the ETERNUS management console and encryption is on a LUN by LUN basis.

There are two primary encryption algorithms supported:

- Fujitsu Original Encryption mode
- AES 128-bit (Advanced Encryption Standard) mode

All data within the system is encrypted and decrypted within the controller as shown below in Figure 4. When encryption is enabled, data stored on disks is all encrypted with the same key. By keeping the key separate from the drives, Fujitsu provides a high level of security if drives are lost or stolen.

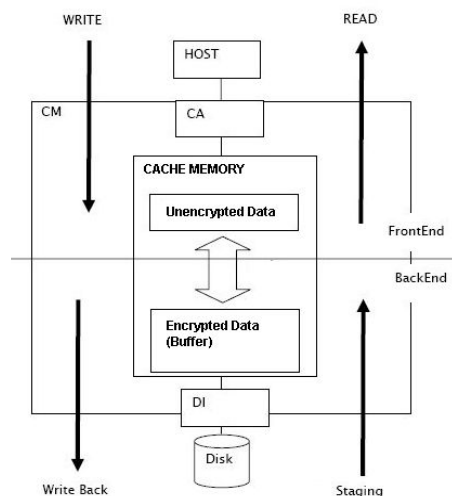


Figure 4: Disk Data Encryption Detail (Source: Fujitsu)



Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

Evaluator Group Comments: *The encryption capabilities of the Fujitsu ETERNUS DX400 S2 can be of significant importance where security is required.*

Additionally, the ETERNUS DX400 S2 supports encryption during replication using standard IPsec security (AES encryption with HMAC-MD5 authentication). This feature is enabled through the use of the optional iSCSI Remote Adapter for remote replication. For more information see the **Remote Replication** section on page 24 in this document.

Mirrored Cache

The controllers each contain cache. Only write data is mirrored from one controller to another controller. With a varying amount of write data, the amount of cache available between the controllers is also variable.

The write data is mirrored between the two controllers over a shared PCI-Express bus link. The bandwidth of the cache on a controller is not specified by Fujitsu. Table 5 below describes the number of controllers and cache capacities for the ETERNUS models.

Battery

ETERNUS DX410 S2 and DX440 S2 protect write data in cache memory that has not yet been written to its intended target drive at the time of a power fault by keeping battery power applied to the controllers long enough so that the write cache data is safely stored in SSD on each of the controllers. Later, after the power fault has been cleared and power has been reapplied to the subsystem, the write data is restored to cache memory and the write data is then correctly written to its original target disk drives.

Battery life is specified at two and a half to three years. The batteries must be replaced if the expiration date has passed. The expiration date can be read through the ETERNUSmgr web-based software.

LUN Management

Fujitsu refers to their LUN masking and LUN management features as 'Host Affinity'. Some vendors term this feature LUN Masking, or LUN Management, or LUN Mapping. LUN access or Host Affinity is controlled by the ETERNUSmgr configuration software, which is included with the system. Host Affinity controls LUN access, which provides a security mechanism at the LUN level. A Host Affinity group is defined as a set of LUNs that are made available to a specific set of servers. The server is identified by the World Wide Name of the host bus adapter.

LUNs can be assigned to multiple servers, with one LUN used in different Host Affinity groups allowing for multi-host data access for cluster or other configurations using host software to coordinate data sharing.

A LUN may be comprised of parts of one or more RAID groups, or a partial RAID group.

Virtual / Thin Provisioning

Fujitsu has added support for virtual capacity of volumes, which is known as "Thin Provisioning." Additionally, Fujitsu supports expansion of a LUN through "Logical Device Expansion" or LDE.



Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

Evaluator Group Comments: *Thin provisioning has become an increasingly important checklist item feature. The effectiveness of thin provisioning in general is application dependent, with some uses such as virtual server and VDI implementations providing some of the best uses.*

RAID Migration

Fujitsu refers to dynamic migration of a LUN as the ability to copy data from one RAID group to another. Since on-line RAID migration copies data from one RAID group to another, the target RAID group must be as large, or larger than the source group. There are few restrictions to RAID migration, allowing administrators to re-architect data layout to support changing requirements. This provides the ability to manage changing capacity, performance or availability requirements for data.

LUN Concatenation

LUN concatenation adds new space for a volume and expands volume capacity that can be used from the server. This feature obtains the area to be used for capacity expansion from unused areas in an existing RAID group, or an unused RAID group. LUN concatenation creates a new volume with the obtained area, and concatenates these volumes to be used as a large capacity volume. LUN concatenation can be used for volume expansion when creating a new volume or expanding an existing volume. Up to 16 volumes may be concatenated. Concatenation of multiple free areas in different or the same RAID group is possible. Concatenation is possible even if the RAID level is different.

Disk Traffic Control (DTC)

DTC is a performance optimization that optimizes the back end disk drive by segmenting sequential and random accesses into separate extents when they both occur simultaneously. The two I/O streams (sequential and random) are processed alternately on a time interval basis to provide optimum usage of the disk drive resource.

Other Features

Fujitsu also supports two modes for specific environments:

- Oracle HARD
- Video Streaming Mode

The Oracle HARD feature provides end-to-end data fidelity. This prevents potential corrupt data from being written to the disk.

The Video streaming mode allows forced prefetching of data when sequential read or write operations are the predominant workload. This can improve performance for video applications.



Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

PERFORMANCE

Performance tests for some past models of the ETERNUS DX400 S2 have been performed and audited by the Storage Performance Council. It is expected that the Model DX440 S2 will complete the SPC-1 and SPC-2 benchmarks. The information will be available at www.storageperformance.org.

Evaluator Group Comments: The benchmark results for the older ETERNUS4000 and DX440 systems are published for both SPC-1 and SPC-2 and show performance levels that match or exceed many of their competitors. The results for the new Model DX440 S2 expected to exceed those of the earlier ETERNUS DX440.



Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

RAS FEATURES

ETERNUS storage systems maintain data integrity and provide data availability via the following mechanisms:

- RAID protection of disk devices
- Redundancy of active components
- Dual power source
- Write cache data de-stage to disk
- Concurrent configuration changes
- Hot spare disk drives
- Concurrent maintenance on components
- Concurrent firmware upgrades
- Parity protection
- ECC (Error Correction Code)
- BlockGUARD

Data Integrity

Data integrity is expected of any storage system and the methods that system vendors employ to ensure data integrity are important for any analysis. Fujitsu has implemented an end-to-end integrity checking mechanism that is called BlockGUARD. BlockGUARD puts an extra eight bytes of information on every 512-byte data block as it enters the system at the channel adapter. These eight bytes are retained with the data as long it is stored within the system. The eight bytes are composed of a four byte cyclical redundancy check (CRC) and four bytes that are the block ID. This information is checked whenever the block is stored on the disk by the device adapter and is also checked when data is retrieved, both at the device adapter and the channel adapter. The BlockGUARD characters are stripped whenever data is sent to the host by the channel adapter.

Data integrity is also a part of the cache function. Data in cache is error-correction code (ECC) protected to be able to detect and correct cache memory errors.

As described previously, Fujitsu also supports Oracle HARD for end-to-end data protection with Oracle databases.

Redundant Copy

When a disk is identified for preventive replacement (before hard failure), the redundant copy function builds the identified disk's data from the other disks in the RAID group to a hot spare. Once that action is complete, the disk at risk is detached and the hot spare is incorporated into the RAID group ensuring data integrity. This is accomplished by using S.M.A.R.T (Self-Monitoring Analysis and Reporting Technology) technology in the disk drive to collect patrol and error statistics.

Redundant Power System

The power system is based on redundant power supplies in the modular enclosures with dual line cord inputs. The modular architecture and dual power inputs ensure the greatest amount of power problem isolation. All power system components are replaceable non-disruptively.



Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

Battery Backup

For ETERNUS DX410 S2 and DX440 S2, the battery backup maintains power for the controller and SSD on the controller. In the event of a power failure, write cache data is saved on the SSD on the controller.

Parity

The PCI and PCI-Express busses have parity generating and checking circuitry that verify hardware integrity for data transfers at the time that they are placed onto and taken off of the busses.

Error Checking and Correction (ECC)

The cache memory on each controller module uses ECC to detect all single and double bit errors and correct single bit errors.

BlockGUARD

The earlier explanation of the data assurance by appending check characters onto every data block is part the error detection system.

Hot Spares

Using hot spare disks is a configuration option. The operation of handling hot spares is as described in the RAID Groups section of this document. After a failed drive is replaced, the data will be copied to the replaced drive and the original spare returned to spare status.

RAID6

With the use of dual parity drives in a RAID6 configuration, the subsystem is better protected in the event of a double disk failure (two disk drives fail concurrently.) RAID6 can be used with any type of drives within the ETERNUS systems.

Concurrent Maintenance

A modular architecture is based on modules being able to be taken out of service while others take over the workload. The ETERNUS DX400 S2 systems operate in a mode where some components can be replaced without failing over to another controller while others require a failover to occur for service. Fujitsu states the following components can have concurrent maintenance to provide access assurance.

- Array Controller, contains – Controllers, cache and channel adapters
- Power Supplies
- Cooling Fans
- Disks

Concurrent Firmware Updates

Firmware updates in the ETERNUS DX400 S2 systems is accomplished by the following operation:

1. Failing over LUN access from the primary controller to the secondary controller
2. Copying the firmware to the controllers flash ROM



Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

3. Then failing the LUNs back their original, primary path
4. Repeating steps 1 – 3 for the second controller module

Remote Call Home

The ETERNUS DX400 S2 systems support remote call home with several remote support features including:

- Send Communication Log
- Automatic sending of system Log
- Manual system log sending

Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

CONNECTIVITY OPTIONS

The following table lists supported servers and operating systems:

Host Vendors	Supported Servers	Supported Operating Systems
Fujitsu	PRIMEQUEST (400/500/500A series)	Windows Server® 2003 for Itanium-based Systems, Windows Server® 2008 for Itanium-based Systems, Red Hat Enterprise Linux AS (V.4 for Itanium), Red Hat Enterprise Linux(V.5 for Itanium)
	PRIMEQUEST (1000 series)	Windows Server® 2003, Windows Server® 2008, Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6, VMware® vSphere 4, VMware vSphere 4.1
	SPARC Enterprise	Solaris 10 Operating System
	PRIMEPOWER	Solaris 8 Operating System, Solaris 9 Operating System, Solaris 10 Operating System
	PRIMERGY	Windows® 2000 Server, Windows Server® 2003, Windows Server® 2003 x 64 Editions, Windows Server® 2008 (32-bit), Windows Server 2008 (64-bit), Red Hat Enterprise Linux AS v.3, Red Hat Enterprise Linux AS v.4, RedHat Enterprise Linux ES v.3, Red Hat Enterprise Linux ES v.4, Red Hat Enterprise Linux 5 (for x86), Red Hat Enterprise Linux 5 (for Intel®64), SUSE Linux Enterprise Server 9 for x86, SUSE Linux Enterprise Server 10 for EM64T, SuSE Linux Enterprise Server 11, x86, SUSE Linux Enterprise Server 11 for EM64T, VMware® vSphere 4, VMware® vSphere 4.1
Sun	Sun Fire, Sun Blade, Sun Enterprise	Solaris 8 Operating System, Solaris 9 Operating System, Solaris 10 Operating System
HP	HP 9000 series	HP-UX 11iv2, HP-UX 11iv3
IBM	pSeries, RS/6000 series	AIX 5.2, AIX 5.3
Others	IA server vendors (various companies)	Windows® 2000 Server, Windows Server® 2003, Windows Server® 2003 x 64 Editions, Windows Server® 2008 (32-bit), Windows Server 2008 (64-bit), Red Hat Enterprise Linux AS v.3, Red Hat Enterprise Linux AS v.4, RedHat Enterprise Linux ES v.3, Red Hat Enterprise Linux ES v.4, Red Hat Enterprise Linux 5 (for x86), Red Hat Enterprise Linux 5 (for Intel®64), SUSE Linux Enterprise Server 9 for x86, SUSE Linux Enterprise Server 9 for EM64T, SuSE Linux Enterprise Server 10, x86, SUSE Linux Enterprise Server 10 for EM64T, VMware® EXS Server 2, VMware ESX Server 3.0 VMware ESX Server3.5

Table 4: ETERNUS DX400 S2 Supported Servers and Operating Systems

Multi-Path IO

Fujitsu provides a multi pathing driver, referred to generically as MPD for the operating platforms supported by the ETERNUS DX400 S2 line. For each platform there is a specific driver package known as ETERNUSmpd.

Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

The ETERNUS DX400 S2 supports either a fixed path access on a per LUN basis, or dynamic multi-path access on a per LUN basis. This feature is different than active/active controllers, and different still from multi-path I/O. The ETERNUS DX400 S2 series permits asymmetric LUN access. This is a mode which Fujitsu refers to as “Assigned CM.”

The ETERNUS DX400 S2 systems assign a controller module (CM) to a specific RA ID Group. The path(s) connecting to the controller normally used will have an active status, while the path(s) connecting to other controllers will have standby status.

Paths connected to the assigned CM are in active status, while paths connected to the non-assigned CM are in standby status. Load-balancing is performed on the paths in active status (called active paths). Paths in standby status (called standby paths) are normally not used for access.

Thus, the ETERNUS DX400 S2 provides active/active controller use with multi-pathing, and asymmetric access to a LUN through a specific array controller (or control module - CM). The diagram below illustrates the configuration for AIX, although this is also the case for other OS environments. Figure 3 below illustrates asymmetric access.

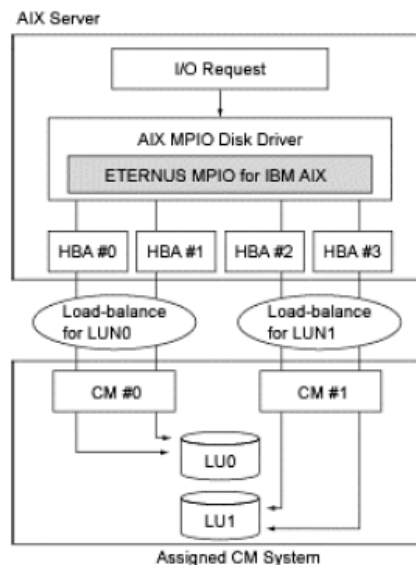


Figure 5: Asymmetric MPIO Access for ETERNUS DX400 S2 (Source: Fujitsu)

Solaris

Fujitsu provides a device driver for Solaris platforms supporting MPIO, utilizing Fujitsu’s multi-path driver, or MPD.

Windows

Fujitsu provides a device driver for Windows platforms supporting MPIO, utilizing Fujitsu’s multi-path driver, or MPD.

Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

Linux

Fujitsu provides a device driver for Linux platforms supporting MPIO, utilizing Fujitsu's multi-path driver, or Linux DM driver.

IBM AIX

For the AIX platform, Fujitsu's MPD is an interface to AIX's native multi-pathing I/O capabilities. The ETERNUS MPIO for AIX is an IBM Path Control Module (PCM) that utilizes AIX's native multi pathing capabilities, through the MPIO interface. ETERNUS systems utilizing this software may be accessed as an MPIO device by AIX. This allows active/active asymmetric logical unit access via AIX.

ETERNUS MPIO for IBM AIX performs load-balancing over the usable paths to improve the access performance.

iSCSI Connectivity

The following table lists supported servers and operating systems for iSCSI host interface.

Supported Servers	Supported Operating Systems
Industry Standard Server	Windows 2000 Server, Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 x64 Editions, Windows Server 2003 R2 x64 Editions, Microsoft Windows Server 2008, Red Hat Enterprise Linux 5, SUSE Linux Enterprise Server 9, SUSE Linux Enterprise Server 10, VMware ESX/ESXi 3.x.x

Table 5: ETERNUS DX400 S2 Supported Servers and Operating Systems – iSCSI

iSCSI Connectivity

The ETERNUS DX400 S2 systems support iSCSI connectivity. There are several options for iSCSI which may be implemented. The following are supported by the ETERNUS DX400 S2 systems:

- Jumbo frames are supported
- iSNS (internet Storage Name Services)
- Changing TCP/IP settings
- CHAP authentication
- Host affinity mode
- Cancel reservation
- Header Digest
- Data Digest

Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

Specification	Model DX410 S2	Model DX440 S2
Number of iSCSI Hosts	1024	1024
Number of Affinity Groups	1024	1024
Affinity Groups per Port	256	256

Table 6: ETERNUS DX400 S2 iSCSI Host Limitations

SOFTWARE FEATURES AND FUNCTIONS

The ETERNUS SF software suite includes a number of offerings that provide options for system and SAN management, data protection and disaster recovery, backup and archiving, local replication and path failover.

Storage Management

SMI-S Interface

The SNIA along with the DTMF have created a set of storage management interface specifications, known as “SMI-S” interface.

Evaluator Group Comment: Prior models of the ETERNUS DX400 S2 have supported and been certified to be interoperable with SMI-S version 1.2 frameworks. Check with the vendor to confirm the status of current models.

ETERNUSmgr (Included)

ETERNUSmgr is a device management application that is included with ETERNUS storage systems, including ETERNUS DX400 S2 systems. The backend program is embedded in the ETERNUS systems and is accessed by a web browser, the ETERNUSmgr frontend software or other operation management software. This browser based management application is used to configure, manage and monitor (including email notification) one or more ETERNUS systems.

ETERNUSmgr Aid (Optional)

ETERNUSmgr Aid adds management value to ETERNUSmgr with its simplified GUI interface and additional management features. Customers will find it easier to manage their ETERNUS storage systems. ETERNUSmgr Aid is optional software that is installed on an administrative server.

ETERNUS SF Storage Cruiser (Optional)

Storage Cruiser is a SAN management application that provides configuration, specification, and monitoring of a SAN environment. It may also be used to manage the relationship between the storage as viewed from the host server and the physical storage, allowing the whole storage system’s configuration and failure conditions to be managed. It will manage the ETERNUS family and also supports other selected vendors products.

ETERNUS SF Storage Cruiser consists of two modules, with an optional third module:

- The Agent feature is installed on each server node being managed
- The Manager feature is installed on the selected administrative server



Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

- (Optional) A remote Client may be installed on any computer used to perform management functions

ETERNUS SF AdvancedCopy Manager (Optional)

The ETERNUS SF AdvancedCopy Manager provides backup and replication with flexible options and works in conjunction with the optional Advance Copy functions of ETERNUS systems. It allows non-stop high-speed backup of major databases such as Oracle and SQL Server. The methods used are Equivalent Copy (EC), One-Point Copy (OPC), differential backup/replication (QuickOPC) and SnapOPC.

Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

Data Protection Software

Fujitsu has termed their point in time copy feature for the ETERNUS series, One Point Copy (OPC). There is local mirroring (Equivalent Copy), along with three different versions of One Point Copy:

- One Point Copy (OPC)
- Quick OPC
- Snap OPC

	Equivalent Copy	One Point Copy (OPC)		
		One Point Copy	QuickOPC	SnapOPC
Copy Type	Mirror	Clone	Clone	PIT
Architecture	Full Copy	Bit Map	Bit Map	Bit Map
Max # Copies	32/volume Maximum concurrent EC, OPC, QOPC or SnapOPC copy sessions: 8192			
Re-sync	Yes	No	Yes	No
Potential Data Exposure	Synchronous—none	Exposure is back to time OPC was activated	Exposure is back to time QuickOPC was activated	Exposure is back to time when snap was taken
Availability for Restore	Mirror available after completion of initial mirror	Available for restore immediately, or able to replace volume after copy complete	Available for restore, or able to replace volume immediately	Available for restore only after command completes
DR, Fail-over Capability	Yes On Mirror suspend or break	Yes Able to replace volume after copy complete	Yes Able to replace volume as soon as command completes	No Not for DR or fail-over, depends on source data
Default Action	Mirror. Copy is a true mirror until it is split	Copy when command issued	Copy when command issued	Copy on Write
Additional Capacity	100%	100%	100%	Changes only
Access to Copy	R/W after split	R/W immediate	R/W immediate	R/O
Performance Impact	Some impact	Small Impact	Minimal	Minimal

Table 7: Fujitsu ETERNUS DX400 S2 Data Protection Options

Note: The maximum number of copy sessions is based on cache available for management of the copy sessions or a maximum of 8192. Cache usage is determined by the size of the volumes being replicated, the number of volumes being replicated and the resolution of the bitmap. The maximum total copy sessions or 32 copy sessions per volume are not available for all configurations.

Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

One Point Copy (OPC)

The copy is a “clone” copy where an exact original copy is made to another physical area. Entire volumes are copied using OPC.

OPC has two copy phases; the logical phase and the physical phase. The logical phase starts after the OPC command is initiated. A bitmap that represents the copy is created in cache memory. That process takes a few seconds and after completion, the copy is available for use. The physical phase then starts, which is the process of the storage system copying the data to the copy location. The physical copy is completed in the background. During the physical phase, host access to data is available. The copied volume is available almost instantly after OPC initiates, upon completion of the logical phase. This is accomplished as follows:

- Read to Target – If the requested block has been previously copied to the target, a normal read is done to the target volume. If the requested block has not been copied, the system will copy the block from the source to the target and then honor the read from the target volume.
- Write to Source – If the requested block has been copied to the target, a normal write operation is done to the source. If the requested block has not been copied yet, the original block is copied to the target and the write is then done to the source.
- Write to Target – If the requested block has been copied to the target, a normal write operation is done. If the block has not been copied, the bitmap is changed to indicate the copy has been done and the normal write is done.

The process of the physical copying persists through a power down or reboot.

There are three speed settings (pacing parameters) to control the amount of storage system resources consumed for the copy operation. The speed or priority settings are controlled via the management software, ETERNUSmgr. The three speed options are **low**, **high** and **automatic**. The automatic setting changes the priority dynamically in response to the system workload.

One Point Copy consumes cache memory for the bitmap and the additional data being copied.

QuickOPC

The first copy is the same as One Point Copy, except a bitmap is used to track all changed blocks from the time the logical copy completed. The design leverages the use of a bitmap, which tracks changed blocks for the volume. Subsequent copies are optimized by only copying the blocks changed since the previous copy.

SnapOPC

SnapOPC operates similar to other Copy on Write point in time copy operations. Rather than copying the new data to the copy volume, the old data is copied to the duplicate volume. Space is saved by only saving the original blocks of data, when new data is written to the SnapOPC volume. By copying only a portion of the source volume, it minimizes copy volume capacity.

Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

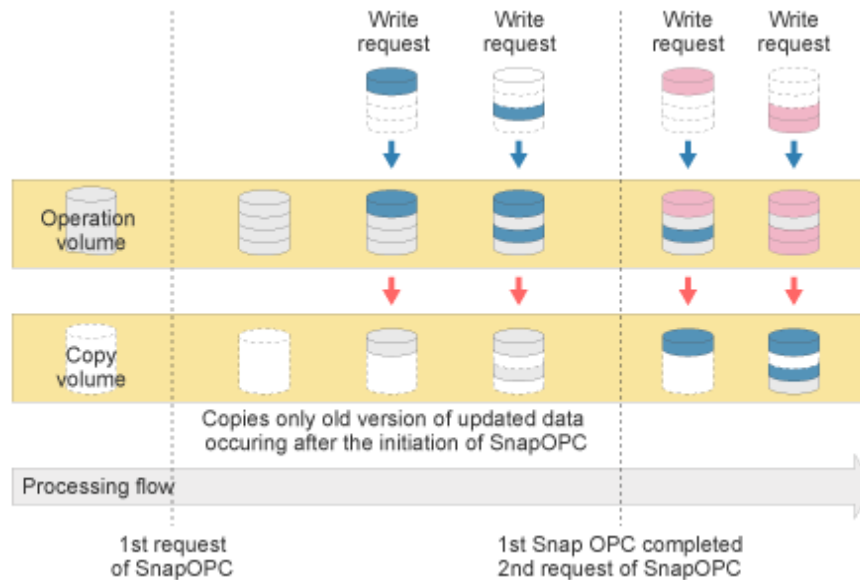


Figure 6: Example of SnapOPC for ETERNUS DX400 S2

SnapOPC+

Similar to SnapOPC, the SnapOPC+ feature backs up the previous data areas to a copy destination area. As a result multiple generations of data may be preserved.

Equivalent Copy (EC)

Fujitsu local mirroring and replication product is known as Equivalent Copy (EC). The mirror is an exact image copy made to another physical area. Entire volumes are copied as initiated by the Fujitsu management software.

The mirror volume must be a local volume, thus the ETERNUS DX400 S2 supports local, in-system mirroring. Like the One Point Copy feature, space for the copy must be reserved previously and be at least the same size. The RAID configuration may be different for the source and target.

The copied volume is not available until the initial copy is complete and either a “detach” (split or break) or a “suspend” command has been issued. During suspend, updates to source and target are tracked. A “resume” command causes resynchronization. Resynchronization copies the incremental updates from the source and overwrites any changed made to the target. The target volume is not available during a resume. If a “detach” has been done, the target volume becomes independent at that time.

The mirror persists through a power down or reboot. As in One Point Copy, Equivalent Copy consumes cache memory for the bitmap and the additional data being copied (and the additional data being copied). To use EC, at least 1GB of cache for each controller module must be installed. While very similar to One Point Copy, the differences are that Equivalent Copy supports the Suspend/Resume function and the mirror volume is not available until the copy is complete.

Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

Remote Replication

Remote Equivalent Copy

Fujitsu's remote replication for the ETERNUS DX400 S2 series is known as Remote Equivalent Copy (REC), which supports replication between multiple ETERNUS DX400 S2s or other ETERNUS storage systems. When referencing remote mirroring specifically, this feature is also known as Remote Equivalent Copy (REC). Fujitsu's RAC product supports both synchronous and asynchronous remote replication.

Remote replication is supported over both FC and IP connections, utilizing the FCP or iSCSI protocols respectively. Utilizing iSCSI replication provides a cost effective method that utilizes TCP/IP connections without requiring an FCIP converter as is typically the case when using FC protocols.

Fujitsu ETERNUS Remote Equivalent Copy provides the following features:

- Maintain an exact byte-for-byte copy of production data in a remote location
- Write order consistency is supported for both Sync and Asynchronous replication
- Consistency groups are supported
- Failover to secondary sites for disaster recovery
- Integration with OPC for secondary site backup operations
- Management of replication pairs using ETERNUS Advanced Copy Management
- Choose between synchronous and asynchronous mirroring for ETERNUS systems

Figure 7 below illustrates the basic example of an asynchronous replication scenario between two ETERNUS DX400 S2 storage subsystems.

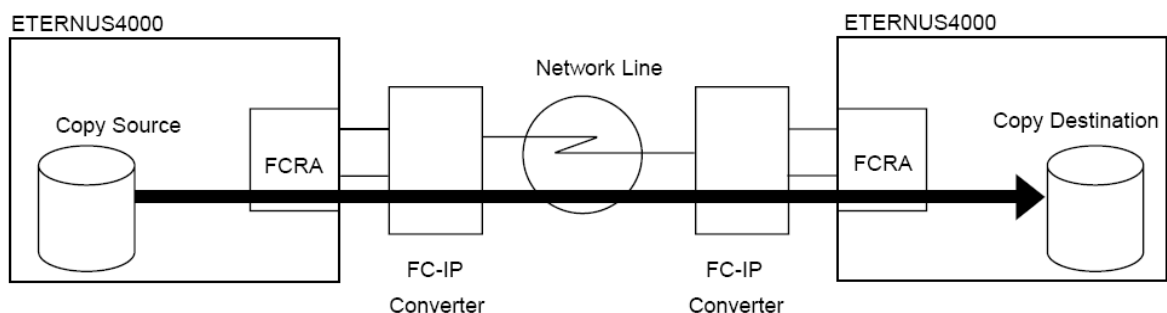


Figure 7: Basic Example of Asynchronous Replication Between Two Subsystems



Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

iSCSI Remote Adapter (optional hardware)

The iSCSI Remote Adapter (iSCSI-RA) provides a number of features for the ETERNUS DX400 S2 and other ETERNUS storage systems.

- Optimized TCP acceleration
 1. Provides optimization for all TCP protocols (FTP, iSCSI, etc)
 2. High bandwidth using large TCP window size
 1. Supports up to 300 Mbps effective TCP transfer rate per card
 2. Supports 500 ms RTT TCP connections
- Encrypted TCP Communication
 3. Utilizes AES encryption with HMAC-MD5 authentication
 4. Supports standard IPsec security
- Bandwidth Control
 5. Traffic shaping and QoS
 6. Automatic TCP window size adjustment with 50 micro-second precision
- Embedded WAN acceleration
 7. All functions supported on one chip on an AMC (Advanced Mezzanine Card)

Evaluator Group Comments: The Fujitsu iSCSI-RA cards provide a significant advantage for customers looking to deploy a cost effective remote replication setup. The use of iSCSI leverages TCP/IP, thereby removing the need for FCIP converters commonly found with FC based replication. Additionally, the encryption provides a level of security mandated by many companies either for security requirements or best practices guidelines. Finally, the inclusion a TCP accelerator can eliminate the need for a separate WAN accelerator in many cases. As a result, it is possible for customers to save well over \$100,000 using the Fujitsu iSCSI-RA cards rather than external encryption and WAN accelerators.

Analysis of Fujitsu ETERNUS DX400 S2
Models DX410 and DX440

SPECIFICATIONS

Model	Width(cm)	Depth(cm)	Height(cm)	Weight(kg)
DX400 S2 Controller Enclosure (rack mount)	48.2	70.0	13.2 (3RU)	54
DX400 S2 2.5" Drive Enclosure (rack mount)	48.2	54.0	8.8 (2RU)	35
DX400 S2 3.5" Drive Enclosure (rack mount)	48.2	55.5	8.8 (2RU)	35

Table 8: ETERNUS DX400 S2 Physical Specifications

Specifications	DX410 S2	DX440 S2
Power – Voltage Requirements	AC 200 – 240V	AC 200 – 240V
No. of Phases	Single	Single
Frequency	50 / 60 Hz	50 / 60 Hz
Max. Power Consumption(W)	11,870	23,500
Max. Heat Dissipation (KJ/h)	43,700	86,600
Temperature Operating	10-40 degrees Centigrade	10-40 degrees Centigrade
Humidity Operating	20-80%RH	20-80%RH

Table 9: ETERNUS DX400 S2 Electrical and Environmental Specifications



**Analysis of Fujitsu ETERNUS DX400 S2
Models DX410 and DX440**

ANNOUNCEMENT HISTORY

Date	Announcement
January, 2009	Announces the new ETERNUS 400DX S2 model400 and model600
April, 2007	Announces support for connectivity into NetApp V-Series virtualization servers.
April, 2006	Announcements of ETERNUS 400DX S2 family of storage systems: Models 80, 100, 300, and 500

Table 10: ETERNUS DX400 S2 Announcement History



Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

EVALUATOR GROUP COMMENTS

The ETERNUS DX400 S2 product line is a solid offering by Fujitsu in the highly competitive midrange storage market. The DX400 S2 series brings many features of the higher end ETERNUS DX8000 down to the midrange, and also shares a common set of management, data protection and replication software with the larger system. The ETERNUS SF Storage Cruiser provides a fully functional management application for the midrange system, which is largely on par with its competitors.

Fujitsu has continued to enhance its platform and server based software features and interoperability. The ETERNUS DX400 S2 is a good choice for enterprises deploying storage for several applications, or workgroup storage consolidation system that provides a wide variety of data protection options and an excellent warranty.

While the Fujitsu ETERNUS line is not known for being first to market with new features, they have built their reputation offering well refined products that are both reliable and offer excellent price performance value. The ETERNUS DX400 S2 continues this tradition, delivering one of the best products in its class with the least amount of recognition. It easily competes with midrange systems from EMC, Hitachi, IBM and NetApp, providing advantages over many of them in several cases.

Strengths:

The ETERNUS DX400 S2 products are well architected storage systems. The features offered are on par with that of most major competitors, and add several industry leading features. The recent addition of thin provisioning complement existing features such multiple RAID levels, including RAID6, along with dynamic RAID migration. These features coupled with excellent performance, expandability and a wide range of data protection options coupled with the industries best in class warranty provide a compelling package.

The ECO Mode feature leverages the idea of spinning down inactive disks, an innovation known as MAID, pioneered by Copan Systems. Fujitsu is one of the first vendors to embrace this idea and incorporate it into their mainstream storage array products. Although the settings are cursory and operate at the RAID group, the capability to selectively spin down drives is an advantage for many deployment scenarios where near-line access capabilities are required, and power and cooling costs are a concern.

The ETERNUS DX400 S2 line offers a wide variety of RAID configuration choices, including RAID6 which is often the best option for large capacity drives. Choice in configuration allows administrators the ability to optimally match RAID configurations to their business use.

Data protection is a critical aspect of storage for large and small deployments. The wide variety of controller based data protection options allow administrators and users to configure the system to provide data protection that best matches their business needs.

The performance of the ETERNUS DX400 S2 is also excellent. The SPC-1 and SPC-2 benchmark results of previous generation systems placed it ahead of many of its competitor's price –



Analysis of Fujitsu ETERNUS DX400 S2 Models DX410 and DX440

performance. Based on past performance, Evaluator Group expects the new systems to perform very well under a variety of workloads and use cases.

Finally, the three year warranty is among the industry's best offering for entry level systems. At a time when most competitors are reducing their warranties to 1 or 2 years, the 3 year Fujitsu warranty will provide peace of mind and cost savings over the life of the product.

Perceived Challenges:

There are only a few issues which should provide any pause for potential users of the ETERNUS DX400 S2. As with many items, the lack of a feature may not be applicable for all environments.

Challenges remain for Fujitsu by way of securing big name distributors, or a large OEM contract. Fujitsu continues to build their direct sales channel and add to their sales and SE force around the world – this certainly is a costly effort, but may be paying off for Fujitsu as they are slowly gaining more credibility with users and continuing to build their brand name in the storage market.

In addition to Fujitsu's solid block offerings, they also partner with both NetApp and Symantec in various geographies to deliver a highly effective NAS solution. However, competitors offering similar bundles will present challenges for Fujitsu to show levels of integration and end-to-end management above the level of their competitors.

Finally, an issue that may be problematic for larger environments is the limited number of Fujitsu personnel available to provide consulting services for their product offering. While Fujitsu has a good presence throughout the world and significant storage expertise, their service offerings are not as diverse or broad as their tier 1 competition, including EMC, HP and IBM.

Copyright 2007-2011 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, inconsequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.